

EXPRESS MAIL LABEL NO.:

(EL803200351US)

5

ADAPTIVE BANDWIDTH UTILIZATION OVER FABRIC LINKS

Kenneth Rose, Mick Jacobs, Jatin Batra

10

BACKGROUND OF THE INVENTION

Local switching networks (e.g., a switching network contained within an office building) may include a switching fabric connecting end devices via line cards. The term switching fabric is a term that describes a distributed switching matrix that establishes a circuit through which data may be transmitted between end devices. A switching fabric may include a stored program control that seeks out a suitable combination of time-slots and multiplexed highways to establish the circuit between the end devices. Multiple highways can co-exist simultaneously being exclusive from/to end devices. The term end device may include desktop computers, printers, routers, or other networking equipment, etc.

15

20

FIG. 1a illustrates, in block diagram form, an exemplary local switching network 100. Local switch network 100 includes a switching fabric 102 (e.g., a cross-bar switching fabric) coupled to line cards 104 through 108. Each of the line cards may be coupled to one or more end devices or other networks. FIG. 1a shows line card 104 coupled to end devices 110 through 114 via common bus 140, line card 106 coupled to end devices 116 through 120 via common bus 142, and line card 108 coupled to end devices 122 through 126 via common bus 144.

25

The local switching network 100 shown in FIG. 1a may employ one of many different communication protocols enabling data communication between one or more end devices 110

through 126 via line cards 104 through 108 and switching fabric 102. FIG. 1a will be described with reference to a communications protocol in which end devices communicate by transferring variable sized data frames or packets to each other via line cards and the switching fabric. Each data packet includes one or more lines of data. The lines of the data packet may be transmitted during consecutive data transmission cycles of the link between the line card and the switching fabric. Lines of the data packet may also be transmitted during non-consecutive transmission cycles of the link between the line card and the switching fabric. Line cards 104 through 108 are coupled to switching fabric 102 via one or more data links. In FIG. 1a, line card 104 is coupled switching fabric 102 via downlink and uplink 128 and 130, respectively, line card 106 is coupled switching fabric 102 via downlink and uplink 132 and 134, respectively, and line card 108 is coupled switching fabric 102 via downlink and uplink 136 and 138, respectively.

FIG. 1b illustrates in block diagram form, relevant portions of a line card 104. As shown in FIG. 1b, line card 104 includes a receiving data buffer 150, a formatter 152, a control circuit 154, and a multiplexer 156. Multiplexer 156 selectively couples an output of formatter 152 and common bus 140 to uplink 130 in response to a control signal generated by control circuit 154. Formatter generates control codes to be transmitted to switching fabric 102. Control circuit couples the output of formatter 152 to uplink 132 in response to formatter generating a control code. Otherwise, control circuit 154 couples common bus 140 to uplink 130 thereby allowing data to be transmitted to switching fabric 102. It is noted lines of data received in line cards are optionally modified by a separate formatter (not shown) before being sent through the multiplexer 156 to fabric 102. For example, the first line of each packet may be modified before passing through the multiplexer 156.

Data buffer 150 typically is a first in first out (FIFO) buffer. Buffer 150 is configured to receive and store lines of data received from switching fabric 102 via downlink 128. The data is stored within buffer 150 until it is transmitted to one of the end devices 110 through 114 (not shown in FIG. 1b) via common bus 140. Common bus 140 may take form in a bi-directional bus. In this embodiment, a switch may be provided (not shown) for selectively coupling the data buffer 150 or the multiplexer 156 to the bi-directional bus. In the alternative, common bus 140 may take form in a pair of unidirectional buses respectively coupled to buffer 150 and multiplexer 156.

Formatter 152 provides a plurality of functions. One function is to monitor buffer 150. Often times, buffer 150 cannot output data to end devices 110 through 114 as fast as data is input to buffer 150 from switching fabric. In this situation, buffer 150 may completely fill as new data is sent from switching fabric 102 for storage therein. Because the buffer has no room for storing the new data, the new data may be lost.

Formatter 152 monitors the quantity of data stored in buffer 150 on an ongoing basis. To this end, formatter 152 receives information from data buffer 150 which identifies the quantity of data stored therein. Formatter compares this data against a single programmable threshold value (e.g., 80% of the full storage capacity of buffer 150). If the quantity of data stored in buffer 150 exceeds this threshold value, formatter 152 generates a control code instructing the switching fabric 102 to stop transmitting data for storage in buffer. Once generated, the stop-transmit control code is transmitted to switching fabric via multiplexer 156 and uplink 130. The switching fabric 102 stops sending data to line card 104 for storage in buffer 150 in response to receiving the stop-transmit control code, thereby decreasing the chances of buffer 150 overflow. However, it is noted that switching fabric may continue to send new data to line card 104 until the switching fabric receives the stop-transmit control code.

As buffer 150 outputs data to one of the end devices 110 through 114 while the switching fabric is in stop-transmit mode of operation, the quantity of data stored therein will decrease. Eventually, formatter eventually will generate a control code instructing the switching fabric to start transmitting data to line card 104. This start-transmit control code will be generated in response to formatter receiving information from buffer 150 that the quantity of data stored therein is less than the single threshold value. The start-transmit control code is transmitted to switching fabric 102 via multiplexer 156 and uplink 130. When the start-transmit control code is received by the switching fabric 102, switching fabric starts to transmit data at the full rate to line card 104 for storage in buffer 104.

Thus, in the system shown in FIG. 1a and 1b, data transmitted by switching fabric 102 is controlled by a simple stop/start control code received from formatter 152; either data is received within data buffer 150 at the full data transmit rate or no data is received at all.